arXiv:1404.7428v1 [cs.AI] 29 Apr 2014

Analysis of Dialogical Argumentation via Finite State Machines

Anthony Hunter
Department of Computer Science,
University College London,
Gower Street, London WC1E 6BT, UK

Abstract

Dialogical argumentation is an important cognitive activity by which agents exchange arguments and counterarguments as part of some process such as discussion, debate, persuasion and negotiation. Whilst numerous formal systems have been proposed, there is a lack of frameworks for implementing and evaluating these proposals. First-order executable logic has been proposed as a general framework for specifying and analysing dialogical argumentation. In this paper[1], we investigate how we can implement systems for dialogical argumentation using propositional executable logic. Our approach is to present and evaluate an algorithm that generates a finite state machine that reflects a propositional executable logic specification for a dialogical argumentation together with an initial state. We also consider how the finite state machines can be analysed, with the minimax strategy being used as an illustration of the kinds of empirical analysis that can be undertaken.

Introduction

Dialogical argumentation involves agents exchanging arguments in activities such as discussion, debate, persuasion, and negotiation (Besnard and Hunter 2008). Dialogue games are now a common approach to characterizing argumentation-based agent dialogues (e.g. (Amgoud, Maudet, and Parsons 2000; Black and Hunter 2009; Dignum, Dunin-Keplicz, and Verbrugge 2000; Fan and Toni 2011; Hamblin 1971; Mackenzie 1979; McBurney and Parsons 2002; McBurney et al. 2003; Parsons, Wooldridge, and Amgoud 2003; Prakken 2005; Walton and Krabbe 1995)). Dialogue games are normally made up of a set of communicative acts called moves, and a protocol specifying which moves can be made at each step of the dialogue. In order to compare and evaluate dialogical argumentation systems, we proposed in a previous paper that first-order executable logic could be used as common theoretical framework to specify and analyse dialogical argumentation systems (Black and Hunter 2012).

In this paper, we explore the implementation of dialogical argumentation systems in executable logic. For this, we focus on propositional executable logic as a special case, and investigate how a finite state machine (FSM) can be generated as a representation of the possible dialogues that can emanate from an initial state. The FSM is a useful structure for investigating various properties of the dialogue, including conformance to protocols, and application of strategies. We provide empirical results on generating FSMs for dialogical argumentation, and how they can be analysed using the minimax strategy. We demonstrate through preliminary implementation that it is computationally viable to generate the FSMs and to analyse them. This has wider implications in using executable logic for applying dialogical argumentation in practical uncertainty management applications, since we can now empirically investigate the performance of the systems in handling inconsistency in data and knowledge.

[1]This paper has already been published in the Proceedings of the International Conference on Scalable Uncertainty Management (SUM'13), LNCS 8078, Pages 1-14, Springer, 2013.

Propositional executable logic

In this section, we present a propositional version of the executable logic which we will show is amenable to implementation. This is a simplified version of the framework for first-order executable logic in (Black and Hunter 2012).

We assume a set of atoms which we use to form propositional formulae in the usual way using disjunction, conjunction, and negation connectives. We construct modal formulae using the $\boxplus$, $\boxminus$, $\oplus$, and $\ominus$ modal operators. We only allow literals to be in the scope of a modal operator. If α is a literal, then each of $\oplus\alpha$, $\ominus\alpha$, $\boxplus\alpha$, and $\boxminus\alpha$ is an **action unit**. Informally, we describe the meaning of action units as follows: $\oplus\alpha$ means that the action by an agent is to add the literal α to its next private state; $\ominus\alpha$ means that the action by an agent is to delete the literal α from its next private state; $\boxplus\alpha$ means that the action by an agent is to add the literal α to the next public state; and $\boxminus\alpha$ means that the action by an agent is to delete the literal α from the next public state.

We use the action units to form **action formulae** as follows using the disjunction and conjunction connectives: (1) If ϕ is an action unit, then ϕ is an action formula; And (2) If α and β are action formulae, then $\alpha \vee \beta$ and $\alpha \wedge \beta$ are action formulae. Then, we define the action rules as follows: If ϕ is a classical formula and ψ is an action formula then $\phi \Rightarrow \psi$ is an **action rule**. For instance, $\mathtt{b}(\mathtt{a}) \Rightarrow \boxplus\mathtt{c}(\mathtt{a})$ is an action rule (which we might use in an example where $\mathtt{b}$ denotes belief, and $\mathtt{c}$ denotes claim, and $\mathtt{a}$ is some information).

Implicit in the definitions for the language is the fact that we can use it as a meta-language

(Wooldridge, McBurney, and Parsons 2005). For this, the object-language will be represented by terms in this meta-language. For instance, the object-level formula $\mathtt{p(a,b)} \rightarrow \mathtt{q(a,b)}$ can be represented by a term where the object-level literals $\mathtt{p(a,b)}$ and $\mathtt{q(a,b)}$ are represented by constant symbols, and $\rightarrow$ is represented by a function symbol. Then we can form the atom $\mathtt{belief(p(a,b)} \rightarrow \mathtt{q(a,b))}$ where `belief` is a predicate symbol. Note, in general, no special meaning is ascribed the predicate symbols or terms. They are used as in classical logic. Also, the terms and predicates are all ground, and so it is essentially a propositional language.

We use a state-based model of dialogical argumentation with the following definition of an execution state. To simplify the presentation, we restrict consideration in this paper to two agents. An execution represents a finite or infinite sequence of execution states. If the sequence is finite, then t denotes the terminal state, otherwise $t = \infty$.

Definition 1 *An* **execution** e *is a tuple* $e = (s_1, a_1, p, a_2, s_2, t)$*, where for each* $n \in \mathbb{N}$ *where* $0 \leq n \leq t$*,* $s_1(n)$ *is a set of ground literals,* $a_1(n)$ *is a set of ground action units,* $p(n)$ *is a set of ground literals,* $a_2(n)$ *is a set of ground action units,* $s_2(n)$ *is a set of ground literals, and* $t \in \mathbb{N} \cup \{\infty\}$*. For each* $n \in \mathbb{N}$*, if* $0 \leq n \leq t$*, then an* **execution state** *is* $e(n) = (s_1(n), a_1(n), p(n), a_2(n), s_2(n))$ *where* $e(0)$ *is the* **initial state***. We assume* $a_1(0) = a_2(0) = \emptyset$*. We call* $s_1(n)$ *the private state of agent 1 at time* n*,* $a_1(n)$ *the action state of agent 1 at time* n*,* $p(n)$ *the public state at time* n*,* $a_2(n)$ *the action state of agent 2 at time* n*,* $s_2(n)$ *the private state of agent 2 at time* n*.*

In general, there is no restriction on the literals that can appear in the private and public state. The choice depends on the specific dialogical argumentation we want to specify. This flexibility means we can capture diverse kinds of information in the private state about agents by assuming predicate symbols for their own beliefs, objectives, preferences, arguments, etc, and for what they know about other agents. The flexibility also means we can capture diverse information in the public state about moves made, commitments made, etc.

Example 1 *The first 5 steps of an infinite execution where each row in the table is an execution state where* `b` *denotes belief, and* `c` *denotes claim.*

n	$s_1(n)$	$a_1(n)$	$p(n)$	$a_2(n)$	$s_2(n)$
0	$\mathtt{b(a)}$				$\mathtt{b}(\neg\mathtt{a})$
1	$\mathtt{b(a)}$	$\boxplus\mathtt{c(a)}$ $\boxminus\mathtt{c}(\neg\mathtt{a})$			$\mathtt{b}(\neg\mathtt{a})$
2	$\mathtt{b(a)}$		$\mathtt{c(a)}$	$\boxplus\mathtt{c}(\neg\mathtt{a})$ $\boxminus\mathtt{c(a)}$	$\mathtt{b}(\neg\mathtt{a})$ $\mathtt{b}(\neg\mathtt{a})$
3	$\mathtt{b(a)}$	$\boxplus\mathtt{c(a)}$ $\boxminus\mathtt{c}(\neg\mathtt{a})$	$\mathtt{c}(\neg\mathtt{a})$		$\mathtt{b}(\neg\mathtt{a})$
4	$\mathtt{b(a)}$		$\mathtt{c(a)}$	$\boxplus\mathtt{c}(\neg\mathtt{a})$ $\boxminus\mathtt{c(a)}$	$\mathtt{b}(\neg\mathtt{a})$
5	...	...	...	...	...

We define a system in terms of the action rules for each agent, which specify what moves the agent can potentially make based on the current state of the dialogue. In this paper, we assume agents take turns, and at each time point the actions are from the head of just one rule (as defined in the rest of this section).

Definition 2 *A* **system** *is a tuple* $(Rules_x, Initials)$ *where* $Rules_x$ *is the set of action rules for agent* $x \in \{1, 2\}$*, and* $Initials$ *is the set of initial states.*

Given the current state of an execution, the following definition captures which rules are fired. For agent x, these are the rules that have the condition literals satisfied by the current private state $s_x(n)$ and public state $p(n)$. We use classical entailment, denoted $\models$, for satisfaction, but other relations could be used (e.g. Belnap's four valued logic). In order to relate an action state in an execution with an action formula, we require the following definition.

Definition 3 *For an action state* $a_x(n)$*, and an action formula* ϕ*,* $a_x(n)$ **satisfies** ϕ*, denoted* $a_x(n) \mid\sim \phi$*, as follows.*

1. $a_x(n) \mid\sim \alpha$ *iff* $\alpha \in a_x(n)$ *when* α *is an action unit*
2. $a_x(n) \mid\sim \alpha \wedge \beta$ *iff* $a_x(n) \mid\sim \alpha$ *and* $a_x(n) \mid\sim \beta$
3. $a_x(n) \mid\sim \alpha \vee \beta$ *iff* $a_x(n) \mid\sim \alpha$ *or* $a_x(n) \mid\sim \beta$

For an action state $a_x(n)$*, and an action formula* ϕ*,* $a_x(n)$ **minimally satisfies** ϕ*, denoted* $a_x(n) \Vdash \phi$*, iff* $a_x(n) \mid\sim \phi$ *and for all* $X \subset a_x(n)$*,* $X \not\mid\sim \phi$*.*

Example 2 *Consider the execution in Example 1. For agent 1 at n = 1, we have* $a_1(1) \Vdash \boxplus\mathtt{c(a)} \wedge \boxminus\mathtt{c}(\neg\mathtt{a})$*.*

We give two constraints on an execution to ensure that they are well-behaved. The first (propagated) ensures that each subsequent private state (respectively each subsequent public state) is the current private state (respectively current public state) for the agent updated by the actions given in the action state. The second (engaged) ensures that an execution does not have one state with no actions followed immediately by another state with no actions (otherwise the dialogue can lapse) except at the end of the dialogue where neither agent has further actions.

Definition 4 *An execution* $(s_1, a_1, p, a_2, s_2, t)$ *is* **propagated** *iff for all* $x \in \{1, 2\}$*, for all* $n \in \{0, \ldots, t-1\}$*, where* $a(n) = a_1(n) \cup a_2(n)$

1. $s_x(n+1) = (s_x(n) \setminus \{\phi \mid \ominus\phi \in a_x(n)\}) \cup \{\phi \mid \oplus\phi \in a_x(n)\}$
2. $p(n+1) = (p(n) \setminus \{\phi \mid \boxminus\phi \in a(n)\}) \cup \{\phi \mid \boxplus\phi \in a(n)\}$

Definition 5 *Let* $e = (s_1, a_1, p, a_2, s_2, t)$ *be an execution and* $a(n) = a_1(n) \cup a_2(n)$*.* e *is* **finitely engaged** *iff (1)* $t \neq \infty$*; (2) for all* $n \in \{1, \ldots, t-2\}$*, if* $a(n) = \emptyset$*, then* $a(n+1) \neq \emptyset$ *(3)* $a(t-1) = \emptyset$*; and (4)* $a(t) = \emptyset$*.* e *is* **infinitely engaged** *iff (1)* $t = \infty$*; and (2) for all* $n \in \mathbb{N}$*, if* $a(n) = \emptyset$*, then* $a(n+1) \neq \emptyset$*.*

The next definition shows how a system provides the initial state of an execution and the actions that can appear in an execution. It also ensures turn taking by the two agents.

Definition 6 *Let* $S = (Rules_x, Initials)$ *be a system and* $e = (s_1, a_1, p, a_2, s_2, t)$ *be an execution.* S **generates** e *iff (1)* e *is propogated; (2)* e *is finitely engaged or infinitely engaged; (3)* $e(0) \in Initials$*; and (4) for all* $m \in \{1, \ldots, t-1\}$

1. *If m is odd, then $a_2(m) = \emptyset$ and either $a_1(m) = \emptyset$ or there is an $\phi \Rightarrow \psi \in Rules_1$ s.t. $s_1(m) \cup p(m) \models \phi$ and $a_1(m) \Vdash \psi$*
2. *If m is even, then $a_1(m) = \emptyset$ and either $a_2(m) = \emptyset$ or there is an $\phi \Rightarrow \psi \in Rules_2$ s.t. $s_1(m) \cup p(m) \models \phi$ and $a_2(m) \Vdash \psi$*

Example 3 *We can obtain the execution in Example 1 with the following rules: (1)* $\mathtt{b(a)} \Rightarrow \boxplus\mathtt{c(a)} \wedge \boxminus\mathtt{c}(\neg\mathtt{a})$*; And (2)* $\mathtt{b}(\neg\mathtt{a}) \Rightarrow \boxplus\mathtt{c}(\neg\mathtt{a}) \wedge \boxminus\mathtt{c(a)}$.

Generation of finite state machines

In (Black and Hunter 2012), we showed that for any executable logic system with a finite set of ground action rules, and an initial state, there is an FSM that consumes exactly the finite execution sequences of the system for that initial state. That result assumes that each agent makes all its possible actions at each step of the execution. Also that result only showed that there exist these FSMs, and did not give any way of obtaining them.

In this paper, we focus on propositional executable logic where the agents take it in turn, and only one head of one action rule is used, and show how we can construct an FSM that represents the set of executions for an initial state for a system. For this, each state is a tuple $(r, s_1(n), p(n), s_2(n))$, and each letter in the alphabet is a tuple $(a_1(n), a_2(n))$, where n is an execution step and r is the agent holding the turn when $n < t$ and r is 0 when $n = t$.

Definition 7 *A finite state machine (FSM)* $M = (States, Trans, Start, Term, Alphabet)$ **represents** *a system* $S = (Rules_x, Initials)$ *for an initial state* $I \in Initials$ *iff*

$(1) States = \{(y, s_1(n), p(n), s_2(n)) \mid$
there is an execution $e = (s_1, a_1, p, a_2, s_2, t)$
s.t. S *generates* e
and $I = (s_1(0), a_1(0), p(0), a_2(0), s_2(0))$
and there is an $n \leq t$
s.t. $y = 0$ *when* $n = t$
and $y = 1$ *when* $n < t$ *and* n *is odd*
and $y = 2$ *when* $n < t$ *and* n *is even* $\}$

$(2) Term = \{(y, s_1(n), p(n), s_2(n)) \in States \mid y = 0\}$

$(3) Alphabet = \{(a_1(n), a_2(n)) \mid$ *there is an* $n \leq t$
and there is an execution e
s.t. S *generates* e
and $e(0) = I$
and $e = (s_1, a_1, p, a_2, s_2, t)\}$

$(4) Start = (1, s_1(0), p(0), s_2(0))$
where $I = (s_1(0), a_1(0), p(0), a_2(0), s_2(0))$

(5)*Trans is the smallest subset of* $States \times Alphabet \times States$ *s.t. for all executions* e *and for all* $n < t$ *there is a transition* $(\sigma_1, \tau, \sigma_2) \in Trans$ *such that*

$$\sigma_1 = (x, s_1(n), p(n), s_2(n))$$
$$\tau = (a_1(n), a_2(n))$$
$$\sigma_2 = (y, s_1(n+1), p(n+1), s_2(n+1))$$

where x is 1 when n is odd, x is 2 when n is even, y is 1 when $n + 1 < t$ and n is odd, y is 2 when $n + 1 < t$ and n is even, and y is 0 when $n + 1 = t$.

Example 4 *Let M be the following FSM where* $\sigma_1 = (\mathtt{1}, \{\mathtt{b(a)}\}, \{\}, \{\mathtt{b}(\neg\mathtt{a})\})$; $\sigma_2 = (\mathtt{2}, \{\mathtt{b(a)}\}, \{\mathtt{c(a)}\}, \{\mathtt{b}(\neg\mathtt{a})\})$; $\sigma_3 = (\mathtt{1}, \{\mathtt{b(a)}\}, \{\mathtt{c}(\neg\mathtt{a})\}, \{\mathtt{b}(\neg\mathtt{a})\})$. $\tau_1 = (\{\boxplus\mathtt{c(a)}, \boxminus\mathtt{c}(\neg\mathtt{a})\}, \emptyset)$; *and* $\tau_2 = (\emptyset, \{\boxplus\mathtt{c}(\neg\mathtt{a}), \boxminus\mathtt{c(a)}\})$. *$M$ represents the system in Ex 1.*

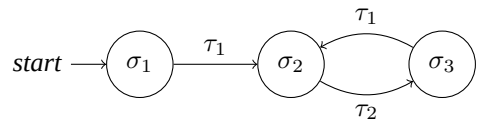

Proposition 1 *For each $S = (Rules_x, Initials)$, then there is an FSM M such that M represents S for an initial state $I \in Initials$.*

Definition 8 *A string ρ* **reflects** *an execution $e = (s_1, a_1, p, a_2, s_2, t)$ iff ρ is the string $\tau_1 \ldots \tau_{t-1}$ and for each $1 \leq n < t$, τ_n is the tuple $(a_1(n), a_2(n))$.*

Proposition 2 *Let $S = (Rules_x, Initials)$ be a system. and let M be an FSM that represents S for $I \in Initials$.*

1. *for all ρ s.t. M accepts ρ, there is an e s.t. S generates e and $e(0) = I$ and ρ reflects e,*
2. *for all finite e s.t. S generates e and $e(0) = I$, then there is a ρ such that M accepts ρ and ρ reflects e.*

So for each initial state for a system, we can obtain an FSM that is a concise representation of the executions of the system for that initial state. In Figure 3, we provide an algorithm for generating these FSMs. We show correctness for the algorithm as follows.

Proposition 3 *Let $S = (Rules_x, Initials)$ be a system and let $I \in Initials$. If M represents S w.r.t. I and* $\mathsf{BuildMachine}(Rules_x, I) = M'$*, then $M = M'$.*

An FSM provides a more efficient representation of all the possible executions than the set of executions for an initial state. For instance, if there is a set of states that appear in some permutation of each of the executions then this can be more compactly represented by an FSM. And if there are infinite sequences, then again this can be more compactly represented by an FSM.

Once we have an FSM of a system with an initial state, we can ask obvious simple questions such as is termination possible, is termination guaranteed, and is one system subsumed by another? So by translating a system into an FSM, we can harness substantial theory and tools for analysing FSMs.

Next we give a couple of very simple examples of FSMs obtained from executable logic. In these examples, we assume that agent 1 is trying to win an argument with agent 2. We assume that agent 1 has a goal. This is represented by the predicate $\mathtt{g(c)}$ in the private state of agent 1 for some argument $\mathtt{c}$. In its private state, each agent has zero or more arguments represented by the predicate $\mathtt{n(c)}$, and zero or more attacks $\mathtt{e(d, c)}$ from $\mathtt{d}$ to $\mathtt{c}$. In the public state, each argument $\mathtt{c}$ is represented by the predicate $\mathtt{a(c)}$. Each agent

can add attacks $\mathtt{e}(\mathtt{d},\mathtt{c})$ to the public state, if the attacked argument is already in the public state (i.e. $\mathtt{a}(\mathtt{c})$ is in the public state), and the agent also has the attacker in its private state (i.e. $\mathtt{n}(\mathtt{d})$ is in the private state). We have encoded the rules so that after an argument has been used as an attacker, it is removed from the private state of the agent so that it does not keep firing the action rule (this is one of a number of ways that we can avoid repetition of moves).

Example 5 *For the following action rules, with the initial state where the private state of agent 1 is* $\{\mathtt{g}(\mathtt{a}),\mathtt{n}(\mathtt{a}),\mathtt{n}(\mathtt{c}),\mathtt{e}(\mathtt{c},\mathtt{b})\}$*, the public state is empty, and the private state of agent 2 is* $\{\mathtt{n}(\mathtt{b}),\mathtt{e}(\mathtt{b},\mathtt{a})\}$*), we get the FSM in Figure 1.*

$$\mathtt{g}(\mathtt{a}) \wedge \mathtt{n}(\mathtt{a}) \Rightarrow \boxplus\mathtt{a}(\mathtt{a}) \wedge \ominus\mathtt{n}(\mathtt{a})$$
$$\mathtt{a}(\mathtt{a}) \wedge \mathtt{n}(\mathtt{b}) \wedge \mathtt{e}(\mathtt{b},\mathtt{a}) \Rightarrow \boxplus\mathtt{a}(\mathtt{b},\mathtt{a}) \wedge \ominus\mathtt{n}(\mathtt{b})$$
$$\mathtt{a}(\mathtt{b}) \wedge \mathtt{n}(\mathtt{c}) \wedge \mathtt{e}(\mathtt{c},\mathtt{b}) \Rightarrow \boxplus\mathtt{a}(\mathtt{c},\mathtt{b}) \wedge \ominus\mathtt{n}(\mathtt{c})$$

The terminal state therefore contains the following argument graph.

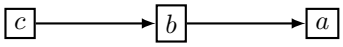

Hence the goal argument $\mathtt{a}$ *is in the grounded extension of the graph (as defined in (Dung 1995)).*

Example 6 *For the following action rules, with the initial state where the private state of agent 1 is* $\{\mathtt{g}(\mathtt{a}),\mathtt{n}(\mathtt{a})\}$*, the public state is empty, and the private state of agent 2 is* $\{\mathtt{n}(\mathtt{b}),\mathtt{n}(\mathtt{c}),\mathtt{e}(\mathtt{b},\mathtt{a}),\mathtt{e}(\mathtt{c},\mathtt{a})\}$*), we get the FSM in Figure 2*

$$\mathtt{g}(\mathtt{a}) \wedge \mathtt{n}(\mathtt{a}) \Rightarrow \boxplus\mathtt{a}(\mathtt{a}) \wedge \ominus\mathtt{n}(\mathtt{a})$$
$$\mathtt{a}(\mathtt{a}) \wedge \mathtt{n}(\mathtt{b}) \wedge \mathtt{e}(\mathtt{b},\mathtt{a}) \Rightarrow \boxplus\mathtt{a}(\mathtt{b},\mathtt{a}) \wedge \ominus\mathtt{n}(\mathtt{b})$$
$$\mathtt{a}(\mathtt{a}) \wedge \mathtt{n}(\mathtt{c}) \wedge \mathtt{e}(\mathtt{c},\mathtt{a}) \Rightarrow \boxplus\mathtt{a}(\mathtt{c},\mathtt{a}) \wedge \ominus\mathtt{n}(\mathtt{c})$$

The terminal state therefore contains the following argument graph.

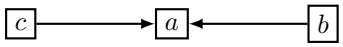

Hence the goal argument $\mathtt{a}$ *is in the grounded extension of the graph.*

In the above examples, we have considered a formalisation of dialogical argumentation where agents exchange abstract arguments and attacks. It is straightforward to formalize other kinds of example to exchange a wider range of moves, richer content (e.g. logical arguments composed of premises and conclusion (Parsons, Wooldridge, and Amgoud 2003)), and richer notions (e.g. value-based argumentation (Bench-Capon 2003)).

Minimax analysis of finite state machines

Minimax analysis is applied to two-person games for deciding which moves to make. We assume two players called MIN and MAX. MAX moves first, and they take turns until the game is over. An **end function** determines when the game is over. Each state where the game has ended is an **end state**. A **utility function** (i.e. a payoff function) gives the outcome of the game (eg chess has win, draw, and loose). The **minimax strategy** is that MAX aims to get to an end state that maximizes its utility regardless of what MIN does

We can apply the minimax strategy to the FSM machines generated for dialogical argumentation as follows: (1) Undertake breadth-first search of the FSM; (2) Stop searching at a node on a branch if the node is an end state according to the end function (note, this is not necessarily a terminal state in the FSM); (3) Apply the utility function to each leaf node n (i.e. to each end state) in the search tree to give the value $value(n)$ of the node; (4) Traverse the tree in post-order, and calculate the value of each non-leaf node as follows where the non-leaf node n is at depth d and with children $\{n_1,..,n_k\}$:

- If d is odd, then $value(n)$ is the maximum of $value(n_1)$,.., $value(n_k)$.
- If d is even, then $value(n)$ is the minimum of $value(n_1)$,.., $value(n_k)$.

There are numerous types of dialogical argumentation that can be modelled using propositional executable logic and analysed using the minimax strategy. Before we discuss some of these options, we consider some simple examples where we assume that the search tree is exhaustive, (so each branch only terminates when it reaches a terminal state in the FSM), and the utility function returns 1 if the goal argument is in the grounded extension of the graph in the terminal state, and returns 0 otherwise.

Example 7 *From the FSM in Example 5, we get the minimax search tree in Figure 5a, and from the FSM in Example 6, we get the minimax search tree in Figure 5b. In each case, the terminal states contains an argument graph in which the goal argument is in the grounded extension of the graph. So each leaf of the minimax tree has a utility of 1, and each non-node has the value 1. Hence, agent 1 is guaranteed to win each dialogue whatever agent 2 does.*

The next example is more interesting from the point of view of using the minimax strategy since agent 1 has a choice of what moves it can make and this can affect whether or not it wins.

Example 8 *In this example, we assume agent 1 has two goals* $\mathtt{a}$ *and* $\mathtt{b}$*, but it can only present arguments for one of them. So if it makes the wrong choice it can loose the game. The executable logic rules are given below and the resulting FSM is given in Figure 4. For the minimax tree (given in Figure 5c) the left branch results in an argument graph in which the goal is not in the grounded extension, whereas the right branch terminates in an argument graph in which the goal is in the grounded extension. By a minimax analysis, agent 1 wins.*

$$\mathtt{g}(\mathtt{a}) \wedge \mathtt{n}(\mathtt{a}) \Rightarrow \boxplus\mathtt{a}(\mathtt{a}) \wedge \ominus\mathtt{n}(\mathtt{a}) \wedge \ominus\mathtt{g}(\mathtt{b})$$
$$\mathtt{g}(\mathtt{b}) \wedge \mathtt{n}(\mathtt{b}) \Rightarrow \boxplus\mathtt{a}(\mathtt{b}) \wedge \ominus\mathtt{n}(\mathtt{b}) \wedge \ominus\mathtt{g}(\mathtt{a})$$
$$\mathtt{a}(\mathtt{a}) \wedge \mathtt{n}(\mathtt{c}) \wedge \mathtt{e}(\mathtt{c},\mathtt{a}) \Rightarrow \boxplus\mathtt{a}(\mathtt{c},\mathtt{a}) \wedge \ominus\mathtt{n}(\mathtt{c})$$

We can use any criterion for identifying the end state. In the above, we have used the **exhaustive end function** giving an end state (i.e. the leaf node in the search tree) which is a terminal state in the FSM followed by two empty transitions. If the branch does not come to a terminal state in the FSM, then it is an infinite branch. We could use a **non-repetitive end function** where the search tree stops when there are no

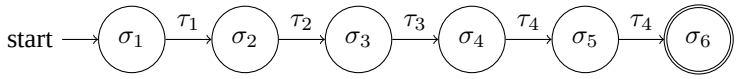

$$\sigma_1 = (1, \{\mathtt{g(a)}, \mathtt{n(a)}, \mathtt{n(c)}, \mathtt{e(c,b)}\}, \{\}, \{\mathtt{n(b)}, \mathtt{e(b,a)}\})$$
$$\sigma_2 = (2, \{\mathtt{g(a)}, \mathtt{n(c)}, \mathtt{e(c,b)}\}, \{\mathtt{a(a)}\}, \{\mathtt{n(b)}, \mathtt{e(b,a)}\})$$
$$\sigma_3 = (1, \{\mathtt{g(a)}, \mathtt{n(c)}, \mathtt{e(c,b)}\}, \{\mathtt{a(a)}, \mathtt{a(b,a)}\}, \{\mathtt{e(b,a)}\})$$
$$\sigma_4 = (2, \{\mathtt{g(a)}, \mathtt{e(c,b)}\}, \{\mathtt{a(a)}, \mathtt{a(b)}, \mathtt{a(c)}, \mathtt{a(c,b)}, \mathtt{a(b,a)}\}, \{\mathtt{e(b,a)}\})$$
$$\sigma_5 = (1, \{\mathtt{g(a)}, \mathtt{e(c,b)}\}, \{\mathtt{a(a)}, \mathtt{a(b)}, \mathtt{a(c)}, \mathtt{a(c,b)}, \mathtt{a(b,a)}\}, \{\mathtt{e(b,a)}\})$$
$$\sigma_6 = (0, \{\mathtt{g(a)}, \mathtt{e(c,b)}\}, \{\mathtt{a(a)}, \mathtt{a(b)}, \mathtt{a(c)}, \mathtt{a(c,b)}, \mathtt{a(b,a)}\}, \{\mathtt{e(b,a)}\})$$

$$\tau_1 = (\{\boxplus\mathtt{a(a)}, \ominus\mathtt{n(a)}\}, \emptyset)$$
$$\tau_2 = (\emptyset, \{\boxplus\mathtt{a(b,a)}, \ominus\mathtt{n(b)}\})$$
$$\tau_3 = (\{\boxplus\mathtt{a(c,b)}, \ominus\mathtt{n(c)}\}, \emptyset)$$
$$\tau_4 = (\emptyset, \emptyset)$$

Figure 1: The FSM for Example 5

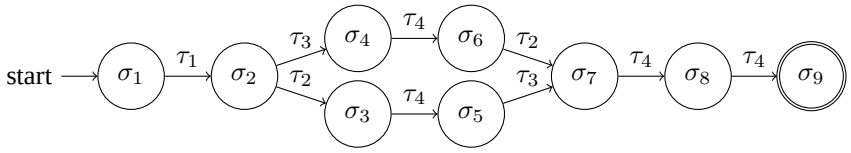

$$\sigma_1 = (1, \{\mathtt{g(a)}, \mathtt{n(a)}\}, \{\}, \{\mathtt{n(b)}, \mathtt{n(c)}, \mathtt{e(b,a)}, \mathtt{e(c,a)}\})$$
$$\sigma_2 = (2, \{\mathtt{g(a)}\}, \{\mathtt{a(a)}\}, \{\mathtt{n(b)}, \mathtt{n(c)}, \mathtt{e(b,a)}, \mathtt{e(c,a)}\})$$
$$\sigma_3 = (1, \{\mathtt{g(a)}\}, \{\mathtt{a(a)}, \mathtt{a(b)}, \mathtt{a(b,a)}\}, \{\mathtt{n(c)}, \mathtt{e(b,a)}, \mathtt{e(c,a)}\})$$
$$\sigma_4 = (1, \{\mathtt{g(a)}\}, \{\mathtt{a(a)}, \mathtt{a(c)}, \mathtt{a(c,a)}\}, \{\mathtt{n(b)}, \mathtt{e(b,a)}, \mathtt{e(c,a)}\})$$
$$\sigma_5 = (2, \{\mathtt{g(a)}\}, \{\mathtt{a(a)}, \mathtt{a(b)}, \mathtt{a(b,a)}\}, \{\mathtt{n(c)}, \mathtt{e(b,a)}, \mathtt{e(c,a)}\})$$
$$\sigma_6 = (2, \{\mathtt{g(a)}\}, \{\mathtt{a(a)}, \mathtt{a(c)}, \mathtt{a(c,a)}\}, \{\mathtt{n(b)}, \mathtt{e(b,a)}, \mathtt{e(c,a)}\})$$
$$\sigma_7 = (1, \{\mathtt{g(a)}\}, \{\mathtt{a(a)}, \mathtt{a(b)}, \mathtt{a(c)}, \mathtt{a(c,a)}, \mathtt{a(b,a)}\}, \{\mathtt{e(b,a)}, \mathtt{e(c,a)}\})$$
$$\sigma_8 = (2, \{\mathtt{g(a)}\}, \{\mathtt{a(a)}, \mathtt{a(b)}, \mathtt{a(c)}, \mathtt{a(c,a)}, \mathtt{a(b,a)}\}, \{\mathtt{e(b,a)}, \mathtt{e(c,a)}\})$$
$$\sigma_9 = (0, \{\mathtt{g(a)}\}, \{\mathtt{a(a)}, \mathtt{a(b)}, \mathtt{a(c)}, \mathtt{a(c,a)}, \mathtt{a(b,a)}\}, \{\mathtt{e(b,a)}, \mathtt{e(c,a)}\})$$

$$\tau_1 = (\{\boxplus\mathtt{a(a)}, \ominus\mathtt{n(a)}\}, \emptyset))$$
$$\tau_2 = (\emptyset, \{\boxplus\mathtt{a(b,a)}, \ominus\mathtt{n(b)})$$
$$\tau_3 = (\emptyset, \{\boxplus\mathtt{a(c,a)}, \ominus\mathtt{n(c)})$$
$$\tau_4 = (\emptyset, \emptyset)$$

Figure 2: The FSM for Example 6

```
01  BuildMachine(Rules_x, I)
02    Start = (1, S_1, P, S_2) where I = (S_1, A_1, P, A_2, S_2)
03    States_1 = NewStates_1 = {Start}
04    States_2 = Trans_1 = Trans_2 = ∅
05    x = 1, y = 2
06    While NewStates_x ≠ ∅
07      NextStates = NextTrans = ∅
08      For (x, S_1, P, S_2) ∈ NewStates_x
09          Fired = {ψ | φ ⇒ ψ ∈ Rules_x and S_x ∪ P ⊨ φ}
10          IfFired == ∅
11          Then NextTrans = NextTrans ∪ {((x, S_1, P, S_2), (∅, ∅), (y, S_1, P, S_2))}
12          Else forA ∈ Disjuncts(Fired)
13            NewS = S_x \ {α | ⊖α ∈ A} ∪ {α | ⊕α ∈ A}
14            NewP = P \ {α | ⊟α ∈ A} ∪ {α | ⊞α ∈ A}
15            Ifx == 1, NextState = (2, NewS, P, S_2) and Label = (A, ∅)
16            Else NextState = (1, S_1, P, NewS) and Label = (∅, A)
17            NextStates = NextStates ∪ {NextState}
18            NextTrans = NextTrans ∪ {((x, S_1, P, S_2), Label, NextState)}
19      If x == 1, then x = 2 and y = 1, else x = 1 and y = 2
20      NewStates_x = NextStates \ States_x
21      States_x = States_x ∪ NextStates
22      Trans_x = Trans_x ∪ NextTrans
23    Close = {σ'' | (σ, τ, σ'), (σ', τ, σ'') ∈ Trans_1 ∪ Trans_2}
24    Trans = MarkTrans(Trans_1 ∪ Trans_2, Close)
25    States = MarkStates(States_1 ∪ States_2, Close)
26    Term = MarkTerm(Close)
27    Alphabet = {τ | (σ, τ, σ') ∈ States}
28    Return (States, Trans, Start, Term, Alphabet)
```

Figure 3: An algorithm for generating an FSM from a system S = $(Rules_x, Initials)$ and an initial state I. The subsidiary function Disjuncts$(Fired)$ is $\{\{\psi^1_1, .., \psi^1_{k_1}\}, .., \{\psi^i_1, .., \psi^1_{k_i}\} \mid ((\psi^1_1 \wedge .. \wedge \psi^1_{k_1}) \vee .. \vee (\psi^i_1 \wedge .. \wedge \psi^1_{k_i})) \in Fired)\}$. For turn-taking, for agent x, $State_x$ is the set of expanded states and $NewStates_x$ is the set of unexpanded states. Lines 02-05 set up the construction with agent 1 being the agent to expand the initial state. At lines 06-18, when it is turn of x, each unexpanded state in $NewStates_x$ is expanded by identifying the fired rules. At lines 10-11, if there are no fired rules, then the empty transition (i.e. $(\emptyset, \emptyset)$) is obtained, otherwise at lines 12-17, each disjunct for each fired rule gives a next state and transition that is added to $NextStates$ and $NextTrans$ accordingly. At lines 19-22, the turn is passed to the other agent, and $NewStates_x$, $States_x$, and $Trans_x$ updated. At line 23, the terminal states are identified from the transitions. At line 24, the MarkTrans function returns the union of the transitions for each agent but for each $\sigma = (x, S1, P, S2) \in Term$, σ is changed to $(0, S1, P, S2)$ in order to mark it as a terminal state in the FSM. At line 25, the MarkStates function returns the union of the states for each agent but for each $\sigma = (x, S1, P, S2) \in Term$, σ is changed to $(0, S1, P, S2)$, and similarly at line 26, MarkTerm function returns the set $Close$ but with each state being of the form $(0, S1, P, S2)$.

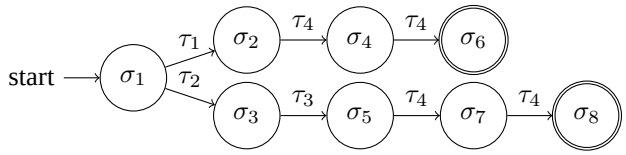

$$\sigma_1 = (1, \{\mathtt{g(a)}, \mathtt{g(b)}, \mathtt{n(a)}, \mathtt{n(b)}\}, \{\}, \{\mathtt{n(c)}, \mathtt{e(c,a)}\})$$
$$\sigma_2 = (2, \{\mathtt{g(a)}, \mathtt{g(b)}, \mathtt{n(a)}\}, \{\mathtt{a(b)}\}, \{\mathtt{n(c)}, \mathtt{e(c,a)}\})$$
$$\sigma_3 = (2, \{\mathtt{g(a)}, \mathtt{g(b)}, \mathtt{n(b)}\}, \{\mathtt{a(a)}\}, \{\mathtt{n(c)}, \mathtt{e(c,a)}\})$$
$$\sigma_4 = (1, \{\mathtt{g(a)}, \mathtt{g(b)}, \mathtt{n(a)}\}, \{\mathtt{a(b)}\}, \{\mathtt{n(c)}, \mathtt{e(c,a)}\})$$
$$\sigma_5 = (1, \{\mathtt{g(a)}, \mathtt{g(b)}, \mathtt{n(b)}\}, \{\mathtt{a(a)}, \mathtt{a(c)}, \mathtt{a(c,a)}\}, \{\mathtt{e(c,a)}\})$$
$$\sigma_6 = (0, \{\mathtt{g(a)}, \mathtt{g(b)}, \mathtt{n(a)}\}, \{\mathtt{a(b)}\}, \{\mathtt{n(c)}, \mathtt{e(c,a)}\})$$
$$\sigma_7 = (2, \{\mathtt{g(a)}, \mathtt{g(b)}, \mathtt{n(b)}\}, \{\mathtt{a(a)}, \mathtt{a(c)}, \mathtt{a(c,a)}\}, \{\mathtt{e(c,a)}\})$$
$$\sigma_8 = (0, \{\mathtt{g(a)}, \mathtt{g(b)}, \mathtt{n(b)}\}, \{\mathtt{a(a)}, \mathtt{a(c)}, \mathtt{a(c,a)}\}, \{\mathtt{e(c,a)}\})$$

$$\tau_1 = (\{\boxplus\mathtt{a(b)}, \ominus\mathtt{n(b)}, \ominus\mathtt{g(a)}\}, \emptyset)$$
$$\tau_2 = (\{\boxplus\mathtt{a(a)}, \ominus\mathtt{n(a)}, \ominus\mathtt{g(b)}\}, \emptyset)$$
$$\tau_3 = (\emptyset, \{\boxplus\mathtt{a(c,a)}, \ominus\mathtt{n(c)}\})$$
$$\tau_4 = (\emptyset, \emptyset)$$

Figure 4: The FSM for Example 8

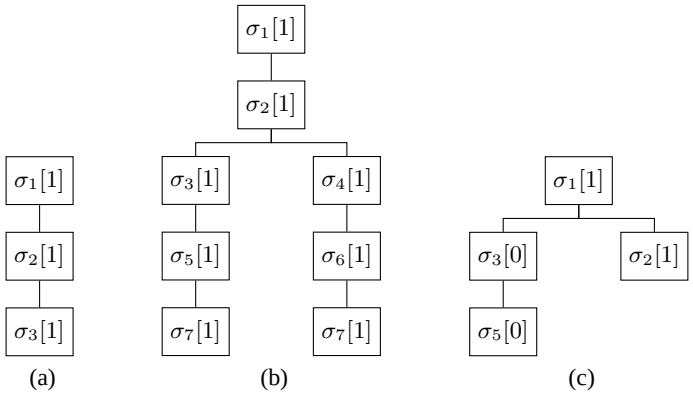

Figure 5: Minimax trees for Examples 7 and 8. Since each terminal state in an FSM is a copy of the previous two states, we save space by not giving these copies in the search tree. The minimax value for a node is given in the square brackets within the node. (a) is for Example 5, (b) is for Example 6 and (c) is for Example 8

new nodes to visit. For instance, for example 4, we could use the non-repetitive end function to give a search tree that contains one branch $\sigma_1, \sigma_2, \sigma_3$ where σ_1 is the root and σ_3 is the leaf. Another simple option is a **fixed-depth end function** which has a specified maximum depth for any branch of the search tree. More advanced options for end functions include **concession end function** when an agent has a loosing position, and it knows that it cannot add anything to change the position, then it concedes.

There is also a range of options for the utility function. In the examples, we have used grounded semantics to determine whether a goal argument is in the grounded extension of the argument graph specified in the terminal public state. A refinement is the **weighted utility function** which weights the utility assigned by the grounded utility function by $1/d$ where d is the depth of the leaf. The aim of this is to favour shorter dialogues. Further definitions for utility functions arise from using other semantics such as preferred or stable semantics and richer formalisms such as valued-based argumentation (Bench-Capon 2003).

Implementation study

In this study, we have implemented three algorithms: The generator algorithm for taking an initial state and a set of action rules for each agent, and outputting the fabricated FSM; A breadth-first search algorithm for taking an FSM and a choice of termination function, and outputting a search tree; And a minimax assignment algorithm for taking a search tree and a choice of utility function, and outputting a minimax tree. These implemented algorithms were used together so that given an initial state and rules for each agent, the overall output was a minimax tree. This could then be used to determine whether or not agent 1 had a winning strategy (given the initial state). The implementation incorporates the exhaustive termination function, and two choices of utility function (grounded and weighted grounded).

The implementation is in Python 2.6 and was run on a Windows XP PC with Intel Core 2 Duo CPU E8500 at 3.16 GHz and 3.25 GB RAM. For the evaluation, we also implemented an algorithm for generating tests inputs. Each test input comprised an initial state, and a set of action rules for each agent. Each initial state involved 20 arguments randomly assigned to the two agents and up to 20 attacks per agent. For each attack in an agent's private state, the attacker is an argument in the agent's private state, and the attacked argument is an argument in the other agent's private state. The results are presented in Table 1.

As can be seen from these results, up to about 15 attacks per agent, the implementation runs in negligible time. However, above 15 attacks per agent, the time did increase markedly, and a substantially minority of these timed out. To indicate the size of the larger FSMs, consider the last line of the table where the runs had an average of 18.02 attacks per agent: For this set, 8 out of 100 runs had 80+ nodes in the FSM. Of these 8 runs, the number of states was between 80 and 163, and the number of transitions was between 223 and 514.

The algorithm is somewhat naive in a number of respects. For instance, the algorithm for finding the grounded extension considers every subset of the set of arguments (i.e. 2^{20} sets). Clearly more efficient algorithms can be developed or calculation subcontracted to a system such as ASPARTIX (Egly, Gaggl, and Woltran 2008). Nonetheless, there are interesting applications where 20 arguments would be a reasonable, and so we have shown that we can analyse such situations successfully using the Minimax strategy, and with some refinement of the algorithms, it is likely that larger FSMs can be constructed and analysed.

Since the main aim was to show that FSMs can be generated and analysed, we only used a simple kind of argumentation dialogue. It is straightforward to develop alternative and more complex scenarios, using the language of propositional executable logic e.g. for capturing beliefs, goals, uncertainty etc, for specifying richer behaviour.

Discussion

In this paper, we have investigated a uniform way of presenting and executing dialogical argumentation systems based on a propositional executable logic. As a result different dialogical argumentation systems can be compared and implemented more easily than before. The implementation is generic in that any action rules and initial states can be used to generate the FSM and properties of them can be identified empirically.

In the examples in this paper, we have assumed that when an agent presents an argument, the only reaction the other agent can have is to present a counterargument (if it has one) from a set that is fixed in advance of the dialogue. Yet when agents argue, one agent can reveal information that can be used by the other agent to create new arguments. We illustrate this in the context of logical arguments. Here, we assume that each argument is a tuple $\langle \Phi, \psi \rangle$ where Φ is a set of formulae that entails a formula ψ. In Figure 6a, we see an argument graph instantiated with logical arguments. Suppose arguments A_1, A_3 and A_4 are presented by agent 1, and arguments A_2, A_5 and A_6 are presented by agent 2. Since agent 1 is being exhaustive in the arguments it presents, agent 2 can get a formula that it can use to create a counterargument. In Figure 6b, agent 1 is selective in the arguments it presents, and as a result, agent 2 lacks a formula in order to construct the counterarguments it needs. We can model this argumentation in propositional executable logic, generate the corresponding FSM, and provide an analysis in terms of minimax strategy that would ensure that agent 1 would provide A_4 and not A_3, thereby ensuring that it behaves more intelligently. We can capture each of these arguments as a proposition and use the minimax strategy in our implementation to obtain the tree in Figure 6b.

General frameworks for dialogue games have been proposed (Maudet and Evrard 1998; McBurney and Parsons 2002). They offer insights on dialogical argumentation systems, but they do not provide sufficient detail to formally analyse or implement specific systems. A more detailed framework, that is based on situation calculus, has been proposed by Brewka (Brewka 2001), though the emphasis is on modelling the protocols for the moves made in dialogical argumentation based on the

Average no. attacks	Average no. FSM nodes	Average no. FSM transitions	Average no. tree nodes	Average run time	Median run time	No. of runs timed out
9.64	6.29	9.59	31.43	0.27	0.18	0
11.47	16.01	39.48	1049.14	6.75	0.18	1
13.29	12.03	27.74	973.84	9.09	0.18	2
14.96	12.50	27.77	668.65	6.41	0.19	13
16.98	19.81	49.96	2229.64	25.09	0.20	19
18.02	19.01	47.81	2992.24	43.43	0.23	30

Table 1: The results from the implementation study. Each row is produced from 100 runs. Each run (i.e. a single initial state and action rules for each agent) was timed. If the time exceeded 100 seconds for the generator algorithm, the run was terminated

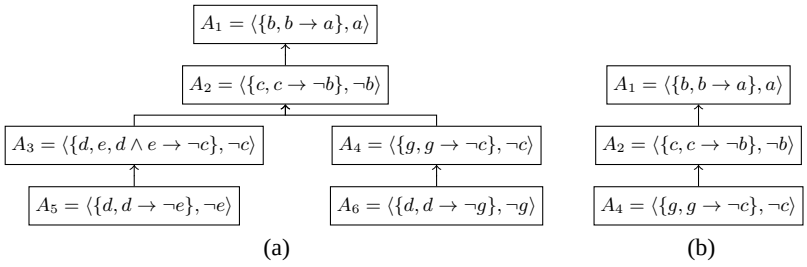
$A_1 = \langle\{b, b \to a\}, a\rangle$
$A_2 = \langle\{c, c \to \neg b\}, \neg b\rangle$
$A_3 = \langle\{d, e, d \wedge e \to \neg c\}, \neg c\rangle$
$A_4 = \langle\{g, g \to \neg c\}, \neg c\rangle$
$A_5 = \langle\{d, d \to \neg e\}, \neg e\rangle$
$A_6 = \langle\{d, d \to \neg g\}, \neg g\rangle$
(a)
$A_1 = \langle\{b, b \to a\}, a\rangle$
$A_2 = \langle\{c, c \to \neg b\}, \neg b\rangle$
$A_4 = \langle\{g, g \to \neg c\}, \neg c\rangle$
(b)

Figure 6: Consider the following knowledgebases for each agent $\Delta_1 = \{b, d, e, g, b \to a, d \wedge e \to \neg c, g \to \neg c\}$ and $\Delta_2 = \{c, c \to \neg b, d \to \neg e, d \to \neg g\}$. (a) Agent 1 is exhaustive in the arguments posited, thereby allowing agent 2 to construct arguments that cause the root to be defeated. (b)Agent is selective in the arguments posited, thereby ensuring that the root is undefeated.

public state rather than on strategies based on the private states of the agents.

The minimax strategy has been considered elsewhere in models of argumentation (such as for determining argument strength (Matt and Toni 2008) and for marking strategies for dialectical trees (Rotstein, Moguillansky, and Simari 2009), for deciding on utterances in a specific dialogical argumentation (Oren and Norman 2009)). However, this paper appears to be the first empirical study of using the minimax strategy in dialogical argumentation.

In future work, we will extend the analytical techniques for imperfect games where only a partial search tree is constructed before the utility function is applied, and extend the representation with weights on transitions (e.g. weights based on tropical semirings to capture probabilistic transitions) to explore the choices of transition based on preference or uncertainty.

References

[Amgoud, Maudet, and Parsons 2000] Amgoud, L.; Maudet, N.; and Parsons, S. 2000. Arguments, dialogue and negotiation. In *European Conf. on Artificial Intelligence (ECAI 2000)*, 338–342. IOS Press.

[Bench-Capon 2003] Bench-Capon, T. 2003. Persuasion in practical argument using value based argumentation frameworks. *Journal of Logic and Computation* 13(3):429–448.

[Besnard and Hunter 2008] Besnard, P., and Hunter, A. 2008. *Elements of Argumentation*. MIT Press.

[Black and Hunter 2009] Black, E., and Hunter, A. 2009. An inquiry dialogue system. *Autonomous Agents and Multi-Agent Systems* 19(2):173–209.

[Black and Hunter 2012] Black, E., and Hunter, A. 2012. Executable logic for dialogical argumentation. In *European Conf. on Artificial Intelligence (ECAI'12)*, 15–20. IOS Press.

[Brewka 2001] Brewka, G. 2001. Dynamic argument systems: A formal model of argumentation processes based on situation calculus. *J. Logic & Comp.* 11(2):257–282.

[Dignum, Dunin-Keplicz, and Verbrugge 2000] Dignum, F.; Dunin-Keplicz, B.; and Verbrugge, R. 2000. Dialogue in team formation. In *Issues in Agent Communication*. Springer. 264–280.

[Dung 1995] Dung, P. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence* 77(2):321–357.

[Egly, Gaggl, and Woltran 2008] Egly, U.; Gaggl, S.; and Woltran, S. 2008. Aspartix: Implementing argumentation frameworks using answer-set programming. In *Proceedings of the Twenty-Fourth International Conference on Logic Programming (ICLP'08),*, volume 5366 of *LNCS*, 734–738. Springer.

[Fan and Toni 2011] Fan, X., and Toni, F. 2011. Assumption-based argumentation dialogues. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI'11)*, 198–203.

[Hamblin 1971] Hamblin, C. 1971. Mathematical models of dialogue. *Theoria* 37:567–583.

[Mackenzie 1979] Mackenzie, J. 1979. Question begging in non-cumulative systems. *Journal of Philosophical Logic* 8:117–133.

[Matt and Toni 2008] Matt, P., and Toni, F. 2008. A game-theoretic measure of argument strength for abstract argumentation. In *Logics in A.I.*, volume 5293 of *LNCS*, 285–297.

[Maudet and Evrard 1998] Maudet, N., and Evrard, F. 1998. A generic framework for dialogue game implementation. In *Proc. 2nd Workshop on Formal Semantics & Pragmatics of Dialogue*, 185198. University of Twente.

[McBurney and Parsons 2002] McBurney, P., and Parsons, S. 2002. Games that agents play: A formal framework for dialogues between autonomous agents. *Journal of Logic, Language and Information* 11:315–334.

[McBurney et al. 2003] McBurney, P.; van Eijk, R.; Parsons, S.; and Amgoud, L. 2003. A dialogue-game protocol for agent purchase negotiations. *Journal of Autonomous Agents and Multi-Agent Systems* 7:235–273.

[Oren and Norman 2009] Oren, N., and Norman, T. 2009. Arguing using opponent models. In *Argumentation in Multi-agent Systems*, volume 6057 of *LNCS*, 160–174.

[Parsons, Wooldridge, and Amgoud 2003] Parsons, S.; Wooldridge, M.; and Amgoud, L. 2003. Properties and complexity of some formal inter-agent dialogues. *J. of Logic and Comp.* 13(3):347–376.

[Prakken 2005] Prakken, H. 2005. Coherence and flexibility in dialogue games for argumentation. *J. of Logic and Comp.* 15(6):1009–1040.

[Rotstein, Moguillansky, and Simari 2009] Rotstein, N.; Moguillansky, M.; and Simari, G. 2009. Dialectical abstract argumentation. In *Proceedings of IJCAI'09*, 898–903.

[Walton and Krabbe 1995] Walton, D., and Krabbe, E. 1995. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. SUNY Press.

[Wooldridge, McBurney, and Parsons 2005] Wooldridge, M.; McBurney, P.; and Parsons, S. 2005. On the meta-logic of arguments. In *Argumentatoin in Multi-agent Systems*, volume 4049 of *LNCS*, 42–56. Springer.